

Sparse point-source removal for full-sky CMB experiments: application to WMAP 9-year data

F.C. Sureau, J.-L. Starck, J. Bobin, P. Paykari, A. Rassat

Laboratoire AIM, UMR CEA-CNRS-Paris 7, Irfu, SAp/SEDI, Service d'Astrophysique, CEA Saclay, F-91191 GIF-SUR-YVETTE Cedex, France

May 22, 2014

Abstract

Missions such as *WMAP* or *Planck* measure full-sky fluctuations of the cosmic microwave background and foregrounds, among which bright compact source emissions cover a significant fraction of the sky. To accurately estimate the diffuse components, the point-source emissions need to be separated from the data, which requires a dedicated processing. We propose a new technique to estimate the flux of the brightest point sources using a morphological separation approach: point sources with known support and shape are separated from diffuse emissions that are assumed to be sparse in the spherical harmonic domain. This approach is compared on both *WMAP* simulations and data with the standard local χ^2 minimization, modelling the background as a low-order polynomial. The proposed approach generally leads to 1) lower biases in flux recovery, 2) an improved root mean-square error of up to 35%, and 3) more robustness to background fluctuations at the scale of the source. The *WMAP* 9-year point-source-subtracted maps are available online.

Key words. Cosmology : CMB, Data Analysis, Methods : Statistical

1. Introduction

Missions such as *Planck* and *WMAP* provide high-resolution full-sky maps of microwave emissions. The primary goal of these missions is to measure the fluctuations in the cosmic microwave background (CMB), the relic radiation of the big bang, and hence provide key information about the birth and evolution of our Universe. These missions also provide full-sky information on the Galactic and extra-Galactic emissions in a frequency range that has not been probed before. One of the main scientific challenges in the CMB data analysis consists of accurately separating the CMB, the various Galactic emissions (synchrotron, free-free, dust and Galactic compact-source emissions, to name a few) as well as the contribution from extra-Galactic sources. In particular, bright compact-source emissions cover a significant fraction of the sky (Planck Collaboration 2011a; Wright et al. 2009), even at high Galactic latitudes, where the CMB is less affected by other diffuse foregrounds. These emissions therefore need to be carefully dealt with for CMB or Galactic data analysis.

From a source-separation perspective, these emissions are difficult to model since the compact sources display both spectral and temporal variability. As reported in Wright et al. (2009), at least one-third of the extra-Galactic *WMAP* sources display temporal variability with high confidence, with more than a 2:1 range in fluxes. On the other hand, extrapolating the fluxes from catalogues obtained at lower frequency than that probed by *WMAP* or *Planck* is erroneous for radio sources (Planck Collaboration 2011b): flat-spectrum emissions from the core of extra-Galactic sources are dominating for Planck or WMAP channels, whereas steep-spectrum lobe emissions dominate at lower frequencies. Finally, their high spectral variability makes it very difficult to estimate their flux with generic source-separation techniques that implicitly rely on the factorization of spatial and spectral information (Tegmark & de Oliveira-Costa 1998; Vielva et al. 2003; López-Caniego et al. 2006). This means that each CMB data set needs to be processed independently to accurately estimate the compact-source fluxes in this data set.

The most distinctive information for separating compact and diffuse emissions is based on morphology, in particular for point sources where the shape of the emission is the point spread function (PSF) of the instrument at the wavelength considered. Consequently, specific pipelines were developed based on this information in the CMB data analysis; the compact sources are first detected and flux-estimated and then are either masked or their contribution subtracted before component separation is performed (Planck Collaboration 2013a).

Numerous approaches based on the morphology of the sources have been proposed to address the initial problem of source detection in the CMB data: morphological image-processing such as SExtractor (Bertin & Arnouts 1996), matched filters (Tegmark & de Oliveira-Costa 1998; Barreiro et al. 2003; López-Caniego et al. 2006), wavelet-based techniques (Cayón et al. 2000; Vielva et al. 2001; Martínez-González et al. 2003; Barreiro et al. 2003; López-Caniego et al. 2006), and Bayesian detection (Hobson & McLachlan 2003; Savage & Oliver 2007; Carvalho et al. 2009; Guglielmetti et al. 2009; Argüeso et al. 2011). Multichannel techniques, which take into account either an estimate of the power spectrum of diffuse emissions (Herranz & Sanz 2008) or the alleged spectral signatures for the point sources (Lanz et al. 2011; Vio et al. 2012) have also been devised. Other multichannel techniques detect the point sources by cancelling the CMB contribution in the data by a weighted sum of the channels (Wright et al. 2009; Ramos et al. 2011; Scodeller et al. 2012). We refer to Paykari & Starck (2012) for a brief review of the techniques. Some compact-source catalogues have been published for the *WMAP* (for instance, Wright et al. 2009; Chen & Wright 2009; Ramos et al. 2011; Scodeller et al. 2012; Bennett et al. 2012) or *Planck* (Planck Collaboration 2011a, 2013b) missions.

When bright sources are detected, their contribution can be removed from the data. Several strategies have been proposed.

- Masking: this is a standard approach where a point-source mask is applied to the data map before it is analysed. This solution is very efficient for CMB power spectrum estimation, and the price to pay is only a small loss of sensitivity at low multipoles (Nolta et al. 2009; Larson et al. 2011). However, for CMB non-Gaussianity analysis the mask may create spurious features in the non-Gaussianity estimators. In addition, masking complicates the Galactic diffuse emission analysis
- Masking and inpainting: to avoid artefacts caused by the mask on large scales, the masked data can be interpolated using inpainting techniques. Several methods have been proposed, such as diffuse inpainting (Zacchei et al. 2011), CMB-constrained realization inpainting (Bucher & Louis 2012; Kim et al. 2012), or sparse inpainting (Abrial et al. 2007; Abrial et al. 2008; Starck et al. 2013b). It was shown that sparse inpainting, based on the sparsity of the CMB in spherical harmonics, does not significantly impact non-Gaussianity measures such as the skewness and kurtosis of the CMB data (Abrial et al. 2007) or the integrated Sachs-Wolfe signal (Dupé et al. 2011) and the weak-lensing signal (Perotto et al. 2010). These inpainting approaches are, however, expected to be inaccurate in estimating the CMB in the masked regions, which would also be detrimental to Galactic studies.
- Fitting: a more progressive approach is to first model the background due to diffuse emissions, then model the shape of the compact sources by modelling the instrument PSF, and finally estimate the flux by aperture photometry or PSF fitting. This is typically performed when the flux and the position are not jointly estimated in the detection algorithm, as performed in the Bayesian detection approaches (where the background can be modelled through a covariance matrix or described by a low-order polynomial). This method was adopted by the *Planck* and *WMAP* teams to estimate the point-source flux (Planck Collaboration 2011a, 2013b; Wright et al. 2009; Bennett et al. 2012), where the background was modelled as a low-order polynomial (baseline or linear background). The few background parameters and the point-source flux were then estimated locally by χ^2 minimization.

Scodeller & Hansen (2012) showed that removing the point-source fluxes instead of masking them provides similar results on the final CMB power spectrum estimation with no biases on higher order statistics. A *cleaned* CMB map is much easier to analyse than a masked CMB map, since there is no need to handle the point-source mask anymore. Another advantage is that it provides a CMB estimate at the point-source positions in the mask (or a point-source-free channel estimate if applied on a channel instead of on the reconstructed CMB map), useful in particular for Galactic studies, while inpainting techniques only try to fill the gaps without destroying the statistical properties of the map.

In this paper, we propose a new approach for point-source removal that is based on a morphological source-separation method, assuming the sources have already been detected. We use a more flexible (and potentially more complex) model for the background to capture its fluctuations more accurately at the scale of the sources. In this approach we assume that the diffuse emissions are sparse in the spherical harmonic domain, while the point sources are sparse in the direct domain, and resolved compact sources are sparse in the wavelet domain. In section 2, we describe this method along with the proposed algorithm, which is adapted from recent convex-optimization techniques that solve the corresponding inverse problem. Results on compact-source removal in full-sky *WMAP* realistic simulations are presented in section 3, where our approach is compared with the standard flux-fitting or low-order polynomial background-fitting as performed by *WMAP* and *Planck* consortia (Planck Collaboration 2011a, 2013b; Wright et al. 2009; Bennett et al. 2012). Sparse point-source removal is then applied on *WMAP* data in Sect. 4, code information is given in Sect. 5, and conclusions and perspectives are drawn in Sect. 6.

2. Sparse point-source removal

2.1. Modelling of the sky

We modelled microwave-channel data over the full sky $y \in \mathbb{R}^{N_e}$ (i.e. the channel map is composed of N_e real-valued pixels) composed of three components $\{x_1, x_2, x_3\} \in \mathbb{R}^{N_e}$: the point sources, the extended compact sources, and the diffuse emissions. In this decomposition, the diffuse background is composed of the CMB itself and the synchrotron, free-free, dust diffuse emissions as well as all the sources with flux below the detection cut. The forward model can be cast as

$$y = x_1 + x_2 + x_3 + n, \quad (1)$$

where $n \in \mathbb{R}^{N_e}$ is an additional noise, assumed to be a real-valued Gaussian random field (but not necessarily stationary). In the associated inverse problem, we need to estimate three unknowns from one equation, which is not possible without additional constraints/information. We now discuss these assumptions and constraints on each of the components.

- Point-source catalogues are available for *WMAP* and *Planck* data sets. Using these catalogues, x_1 can be modelled as a set of Dirac functions $\{\delta_p\}_p$ convolved by the instrumental beam b : $x_1 = \sum_{p=1}^{N_p} f_p \delta_p * b$, where N_p is the number of point sources that need to be removed and f_p is the flux of the p th point source. The source fluxes are expected and enforced to be positive.

- In case of extended compact sources, the morphological information is weaker. Several extended sources were detected and flagged in the *Planck* data set, mostly located close to the Galactic centre, while in the *WMAP* data such sources were not considered (detection was performed outside a Galactic mask). In this work, we assumed that these emissions can be modelled through a sparse decomposition in the undecimated isotropic spherical wavelet dictionary (Starck et al. 2006): $x_2 = \sum_{k_1}^{N_e} \sum_{j=1}^{N_s} w_{jk} \theta_{jk}$, where N_s is the number of wavelet scales, θ_{jk} corresponds to a wavelet atom at scale j and position k , and we have overall $N_w = N_e N_s$ wavelet coefficients at all scales. Assuming sparsity of this Galactic component, only a few wavelet coefficients w_{jk} at selected locations k close to the Galactic centre are assumed significantly away from zero.

- The background emission x_3 is much more complex to model accurately. It is usually represented as a local low-order polynomial background (either baseline, or first order) because of its local smoothness, and in this work it is characterized by its sparsity in the spherical harmonic domain. This assumption of sparsity in CMB applications is motivated in Starck et al. (2013a). x_3 is also assumed band-limited because of the beam effect.

These components can only be separated if the components are sparse enough in mutually incoherent dictionaries, such as a dirac basis and a spherical harmonic basis (Donoho & Huo 2001; Donoho & Elad 2003). Said differently, compact-source emissions are expected to decrease the sparsity level of the data measured on spherical harmonics, which will be the key ingredient driving the separation process in our approach.

2.2. Associated inverse problem

Using all these hypotheses and constraints, we consider the following inverse problem:

$$\underset{f \in \mathcal{C}, w \in \mathbb{R}^{N_w}, a \in \mathcal{D}}{\text{minimize}} \quad \gamma \|a\|_1 + \beta \|w\|_1 \quad \text{s.t.} \quad \|y - (\mathbf{B}f + \mathbf{M}Ww + \mathbf{S}a)\|_{2,\Sigma} < \epsilon, \quad (2)$$

where

- $f \in \mathbb{R}^{N_p}$ is a vector containing the flux of the N_p point sources, \mathcal{C} is the positive orthant of \mathbb{R}^{N_p} to enforce the positivity constraint, and $\mathbf{B} \in \mathbb{R}^{N_e \times N_p}$ is a matrix operating on fluxes and implementing the local projection of the beam at the location of the point sources;
- $w = \{w_{jk}\}_{j,k} \in \mathbb{R}^{N_w}$ is a vector containing the N_w wavelet coefficients at all scales, \mathbf{W} represents the matrix implementing the spherical wavelet transform (containing each wavelet atom as a column), \mathbf{M} corresponds to a binary mask equal to one where compact extended sources are expected. We promote sparsity of x_2 by making use of the ℓ_1 norm of w as a penalization term;
- $a \in \mathbb{C}^{N_l}$ contains the N_l complex spherical harmonic coefficients of x_3 and \mathcal{D} is the space of band-limited signals on the sphere (up to a chosen multipole ℓ_{max} depending on the resolution of the channel). $\mathbf{S} \in \mathbb{C}^{N_e \times N_l}$ is the matrix describing the orthogonal spherical harmonic transform: $x_3 = \mathbf{S}a$. Sparsity of x_3 is enforced through the use of a ℓ_1 norm penalization term on a , which is computed as the sum of the modulus of the complex multipole as in the sparse inpainting described in Starck et al. (2013a)

- Σ is the (non-stationary) noise covariance matrix and $\|x\|_{2,\Sigma}^2 = x^T \Sigma^{-1} x$ denotes the square of the ℓ_2 norm weighted by Σ^{-1} .

The reconstructed point-source-free map \tilde{y} can be obtained by $\tilde{y} = y - \mathbf{B}f$. The problem described in Equation 2 is a convex problem and is related to a constrained morphological component analysis (Starck et al. 2004, 2005) or a basis pursuit denoising problem (Chen & Donoho 1998) with a deconvolution step. The principal advantage of this constrained formulation lies in having only a few hyperparameters to set, which are easily interpretable: the expected noise level ϵ ; the balance between ℓ_1 norm in the wavelet and spherical harmonic domains. More explicitly, the higher the ratio (γ/β) , the sparser the solution for x_3 is in the spherical harmonic domain, and conversely, the lower the ratio, the sparser the estimate of x_2 in the wavelet domain. Setting them equal would lead to penalize both non-sparse solutions for x_2 in the wavelet domain and x_3 in the spherical harmonic domain, as performed in morphological component analysis (Starck et al. 2004, 2005). Note that once the ratio is set, changing their modulus does not change the inverse problem.

The main difficulty here lies in controlling the interplay between sparsity and data fidelity constraints: how can we efficiently estimate a sparse solution (more precisely with minimal ℓ_1 norm), knowing that many combinations of x_1 , x_2 and x_3 can satisfy the data fidelity constraint?

In the next section, we propose an algorithm to solve this problem.

2.3. Proposed algorithm

Inverse problems as described by Equation 2 can be solved using primal-dual approaches (e.g. Chambolle & Pock 2011; Briceño Arias & Combettes 2011; Becker et al. 2011; Combettes & Pesquet 2012). The proposed algorithm was derived from Chambolle & Pock (2011); it requires only one application of the costly spherical harmonic and wavelet transforms and one application of their adjoint per iteration and does not require sub-iterations. The algorithm is set as follows:

Morphological component analysis with a primal-dual approach (SPSR)

- 1- Choose $(a^0, f^0, w^0, t^0) \in \mathbb{C}^{N_t} \times \mathbb{R}^{N_p} \times \mathbb{R}^{N_w} \times \mathbb{R}^{N_e}$ and $\bar{a}^0 = a^0$, $\bar{f}^0 = f^0$, $\bar{w}^0 = w^0$.
Choose also the hyperparameters γ, β and τ, σ s.t. $\tau\sigma < \frac{1}{3}$ (assuming normalized $\|S\|_2 = \|MW\|_2 = \|B\|_2 = 1$).
- 2- Iterate ($n \geq 0$):

$$\left\{ \begin{array}{l} r_d^n = t^n + \sigma \Sigma^{-1/2} (\mathbf{B} \bar{f}^n + \mathbf{S} \bar{a}^n + \mathbf{M} \mathbf{W} \bar{w}^n - y) \\ t^{n+1} = \begin{cases} 0 & \text{if } \|r_d^n\|_2 \leq \sigma \epsilon \\ (1 - \frac{\epsilon \sigma}{\|r_d^n\|_2}) r_d^n & \text{otherwise} \end{cases} \\ a^{n+1} = \mathcal{ST}_{\tau\gamma} (P_{\mathcal{D}}(a^n - \tau \mathbf{S}^\dagger \Sigma^{-1/2 \dagger} t^{n+1})) \\ f^{n+1} = P_{\mathcal{C}} (f^n - \tau \mathbf{B}^\dagger \Sigma^{-1/2 \dagger} t^{n+1}) \\ w^{n+1} = \mathcal{ST}_{\tau\beta} (w^n - \mathbf{W}^\dagger \mathbf{M} \Sigma^{-1/2 \dagger} t^{n+1}) \\ \bar{f}^{n+1} = 2f^{n+1} - f^n \\ \bar{a}^{n+1} = 2a^{n+1} - a^n \\ \bar{w}^{n+1} = 2w^{n+1} - w^n \end{array} \right. \quad (3)$$

where $P_{\mathcal{C}}$ is the projection onto the positive orthant (i.e. sets to 0 negative fluxes) and $P_{\mathcal{D}}$ is the projection onto the set of considered band-limited signals (i.e. sets to 0 all multipoles greater than the chosen ℓ_{max}); $\mathcal{ST}_{\tau\beta}$ is the standard soft-thresholding operator applied component-wise:

$$[\mathcal{ST}_{\tau\beta} x]_i = x_i \left[1 - \frac{\tau\beta}{|x_i|} \right]_+, \quad (4)$$

where $|x_i|$ is the complex modulus of x_i for complex vectors. From Chambolle & Pock (2011), the sequence $((f^n, a^n, w^n), t^n)$ converges to a saddle point of the primal-dual problem with a restricted dual gap decreasing as $\mathcal{O}(1/n)$ (first-order method). We initialised the algorithm with null images.

The sparse-source-removal (SPSR) algorithm has four parameters to set: β , γ , τ , and σ . The last two are the primal and dual steps, respectively, and drive the convergence speed of the algorithm.

The ratio of the first two parameters is the hyperparameter that controls the balance in the sparsity of the decomposition of the extended compact sources and the diffuse background as discussed in the previous section. The choice of their modulus also affects the convergence speed of the algorithm: a high modulus leads to a slowly built sparse approximation, or said differently, the algorithm would provide very sparse solutions at the beginning of the algorithm that are far from satisfying the data fidelity constraint, however; a low modulus, in contrast, leads to building approximations that satisfy the data fidelity constraint early in the iterations, but are not as sparse. In practice, several values need to be experimentally tested to obtain an algorithm with reasonable convergence speed and to derive the required number of iterations.

To assess the relative performance of the proposed approach, flux estimates obtained with SPSR were compared with the low-order polynomial fitting approach (Bennett et al. 2012; Planck Collaboration 2011a, 2013b). Fluxes were estimated in a 3.5σ region as previously recommended in Wright et al. (2009) using a Levenberg-Marquardt algorithm with local χ^2 minimization (using the C++ library ALGLIB (www.alglib.net), Sergey Bochkhanov and Vladimir Bystritsky). The local background was either modelled with a baseline or a first-order polynomial (**FIT-C** and **FIT-L** in the following) and a positivity constraint was applied on the fluxes. Clustered point sources with overlapping fitting regions were jointly fitted for. However, in practice, such situations were scarce in the *WMAP* simulations reported.

3. Results on synthetic *WMAP* simulations

3.1. *Planck Sky Model* simulations at *WMAP* frequencies

The *Planck Sky Model* (PSM) software (Delabrouille et al. 2013) was employed to simulate *WMAP*-like data. Each one of the five *WMAP* frequency channels included the following:

- A diffuse component, comprised of a Gaussian CMB, generated from a six-parameter Λ -CDM model (with default values from *WMAP* 7-year data combined with BAO and the Hubble constant measurements, and with added C_ℓ lensing) and synchrotron, free-free, thermal-dust and spinning-dust emissions, simulated with default PSM parameters.
- A compact component, constituted by radio, infrared, and strong ultra-compact HII region emissions, as well as a far-infrared background (model jgn2005 in the PSM). More specifically, radio sources are essentially derived from observations conducted between 0.85 GHz and 4.85 GHz. Each one of these ~ 2 million simulated sources possessed its own spectrum, composed of power laws with spectral indices varying across pre-defined bands of electromagnetic frequencies. As reported in Delabrouille et al. (2013), this catalogue is expected to faithfully reproduce the clustering properties of the radio sources and the observed/modelled source number counts at 5 and 20 GHz. In the PSM, the infrared sources were derived from the IRAS point-source catalogue and faint-source catalogue, with assumed modified blackbody emissions to extrapolate the flux at lower frequencies. Simulated sources were also added near the Galactic plane and in IRAS gaps to obtain the same mean surface density as a function of flux (down to 80mJy) as in the regions well covered by IRAS. Emissions from ultra-compact HII regions were also derived from IRAS, with flux extrapolated according to a modified blackbody fit with added low-frequency flux with a free-free spectral index when a radio-counterpart is found. Note that compact emissions from the Galaxy were also simulated as part of diffuse emissions. Finally, the simulated far-infrared background is composed of realistic distribution of clustered point sources with assumed flux-dependent spectral indices and reproduced *Planck* observations reasonably well.
- Noise modelled as a non-stationary Gaussian random field, with variance in each pixel derived from the *WMAP* 9-year hit maps; the channel bandpasses were modelled as diracs located at the centre frequency, and the beams were assumed Gaussian with a full-width at half maximum given by the beam size provided by the *WMAP* consortium.

Overall, the PSM simulations reflect the complexity in estimating the strong point-source fluxes (which is essential to test the robustness of the proposed approach), with complex diffuse background, unresolved sources with clustering properties below the detection limit, strong compact emissions in the Galactic plane, and non-stationary noise. A more comprehensive description of the PSM is provided in Delabrouille et al. (2013). A patch extracted from these *WMAP* 9-year simulated data is represented in Figure 1 at each of the five *WMAP* frequency channels to illustrate the complexity of the source separation, because the resolution, noise, and background change in each of the channels. At low-frequency channels, the background fluctuates at the scale of the point sources, and noise contribution to the flux estimates are lower than in the high-frequency channels, where noise becomes important at the scale of the source, as reported in Wright et al. (2009).




Figure 1. Patch of the simulated sky centred on two detected point sources at the *WMAP* wavelengths.

3.2. Flux recovery in noise-limited catalogue

3.2.1. Catalogue generation

First, a catalogue of point sources was independently generated from the simulations for each *WMAP* channel by retaining the sources with flux above 5σ of the estimated radiometer noise (Wright et al. 2009). The number of sources retained for each channel and the flux cut are displayed in Table 1. This scenario reflects an optimistic situation where only noise would be a limiting factor for source detection (and not the background). Consequently, it allows us to investigate how well point sources can be estimated even when their flux is similar to the background levels. Note that the cut in flux increases with frequency (because of the increasing level of the noise) and our sources are mainly radio sources (i.e. with flux decreasing with increasing frequency). Consequently, fewer sources are detected at higher frequency channels.

Table 1. Characteristics of the point-source catalogue generated from simulations.

Channel	Flux Cut (mJy)	Source Number
K	111	1094
Ka	183	1093
Q	217	1037
V	383	683
W	687	330

3.2.2. SPSR working conditions

The working conditions of SPSR were chosen as follows: the noise hyperparameter ϵ corresponds in the following to the 95th percentile of a $\chi^2(N_e)$ distribution according to the whitened noise statistics. For all channels the hyperparameter ratio was set to $\beta/\gamma = 1$ to favour both sparsity of x_2 in the wavelet domain and x_3 in the spherical harmonic domain, as discussed before.

A range of values for the modulus of $\beta = \gamma$ was first tested and turned out to be crucial for the convergence speed of the algorithm. This is illustrated in Figure 2, where SPSR was run with a high number of iterations for channel Ka. Setting the modulus too high (or too low) leads to many iterations necessary to fulfill the data-fidelity constraint (or minimizing the ℓ_1 penalty for components that satisfy the data-fidelity term). For instance, for values of $\beta = 0.001M$ (or $\beta = 0.005M$) the data-fidelity constraint is satisfied in the first one thousand iterations, but the sparsity-penalty term decreases very slowly and has not converged after 50000 iterations. Conversely, $\beta = 0.1M$ (or $\beta = 0.05M$) leads to very sparse approximations that do not satisfy the data-fidelity constraint even after 50000 iterations. We chose the value $\beta = 0.01M$ for all channels, which led to the best convergence properties among the tested values.

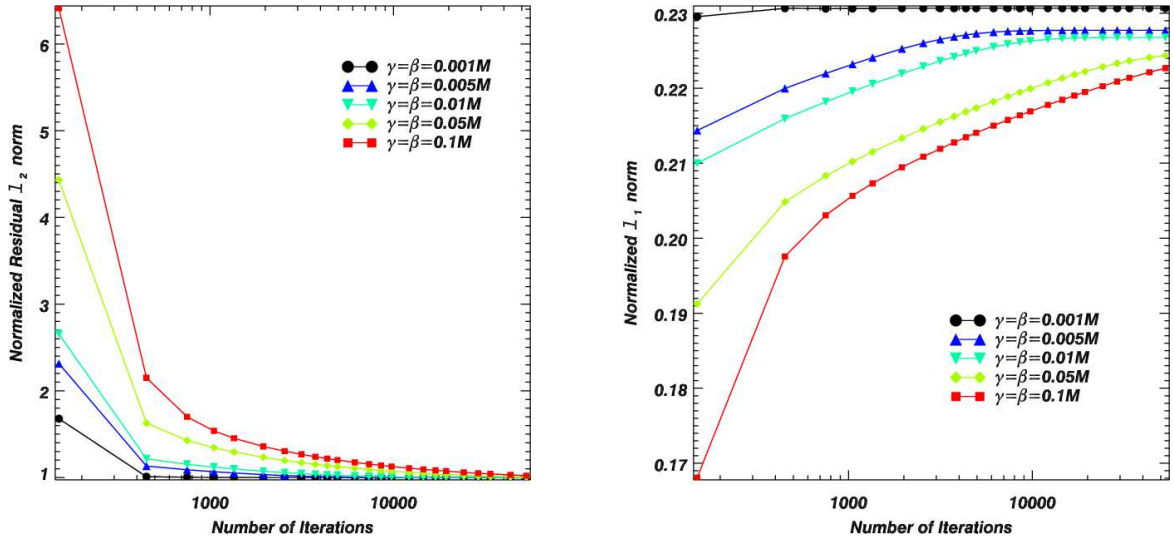


Figure 2. Evolution with the number of iterations of the normalized ℓ_2 norm of the residuals (left) and normalized ℓ_1 norm of the diffuse and compact component for channel Ka and for various values for the modulus of $\beta = \gamma$. Reaching the upper limit of the data-fidelity constraint leads to a normalized residual ℓ_2 norm of 1. The ℓ_1 norm was normalized to the ℓ_1 norm of the input data in direct space (sum of the absolute values of the pixels in the frequency map). M corresponds to the highest modulus of the spherical harmonic coefficients in the input image.

Finally, the number of iterations needs to be set. In addition to the modulus of $\beta = \gamma$, the convergence speed of the algorithm depends on the conditioning of the beam matrix \mathbf{B} . The number of iterations therefore varies from channel to channel, requiring more iterations for lower resolution. The flux estimates at several stages of the proposed algorithm are illustrated in Figure 3 for channels K and W, which correspond to the lowest and highest resolution and to the highest and lowest noise, respectively. In practice, 13350 iterations were chosen for channel K and 9750 iterations for all other channels as a compromise between processing time and level of convergence. We also ensured that the statistics measured were marginally changing with respect to the relative difference between methods when reaching this number of iterations. Generally, the large number of iterations in both cases is caused by the absence of regularity in the inverse problem of Equation 2, which leads to first-order algorithms such as SPSR. To give a practical example of the computational cost of SPSR, 9750 iterations for the channel Ka represent ~ 80 hours of processing in a cluster with 20 Intel Xeon 2 GHz processors, essentially spent in the spherical harmonic and wavelet transforms.




Figure 3. Flux recovery with the proposed approach as a function of the number of iterations. After 13350 (1350) iterations, the flux estimates have converged for channel K (W). Note that the fluxes are displayed in logarithmic scales to visualize the convergence speed for both low- and high-flux sources.

3.2.3. Flux recovery and point-source subtraction

We compared point-source flux recovery in different flux bands for the FIT-C, FIT-L, and SPSR approaches in the simulated channels. The bias, root-mean square errors (RMSE), and mean absolute deviation (MAD) (less sensitive to outliers) are displayed in Table 2. For all methods, the RMSE or MAD varies from channels K to W as a consequence of increasing noise contribution and decreasing background fluctuations at the scale of the point sources, as previously observed in Wright et al. (2009). When the noise contribution to flux uncertainty is predominant (in channel W), FIT-C and FIT-L have generally lower biases, RMSE, and MAD than SPSR. Note, however, that the improvement is very small (at most 6% difference in RMSE, bottom row of Table 2). Therefore, it seems that in the noise-dominating case, fitting a few local parameters for the diffuse background, as in FIT-C and FIT-L, gives slightly more robust flux estimates than SPSR, where many more global parameters are estimated. The situation is the opposite when background fluctuations are the foremost contribution as in the first four channels: SPSR now leads to much lower RMSE and MAD than FIT-C or FIT-L, as well as lower biases (for the first three channels). In particular, point sources with flux lower than 1Jy are better recovered by about 12%, 26%, 32% and 8% in terms of RMSE for channels K to V, with much lower biases for the first three channels (more than 70mJy better for channel K, for instance).

Statistics on the distribution of the error in the different flux bands are displayed in Figure 4. In these box-plots both quartiles (three horizontal bars making the "box") and extreme values (upper and lower horizontal bars) are represented for the various approaches and the different channels. First, the same conclusions can be drawn as in Table 1 by focusing on the interquartile range and median for the different channels: SPSR outperforms FIT-C and FIT-L, except for channel W, where results are only slightly degraded. In particular, Figure 5 illustrates that the bias in the lowest frequency channels is higher than the noise level in the residual maps for FIT-C and FIT-L, but not for SPSR.

These box-plots also illustrate that the proposed approach leads to lower extreme errors, while FIT-L or FIT-C can both fail with large flux errors. An example of such an incorrect flux-fitting is illustrated in Figure 6 for channel K, where FIT-C and FIT-L lead to overestimated flux in a complex background region that cannot be accurately modelled with low-order polynomials. In contrast, SPSR successfully estimates the background and therefore gives a better flux estimate.

3.3. Flux recovery with internal multichannel catalogue derived from simulations

The noise-free catalogue is useful to assess whether the more flexible background model in SPSR indeed better captures the background fluctuation than low-order local polynomials, in particular, to recover low-flux sources. However, in practice detection is performed internally for each channel, and a catalogue is built by combining these multichannel detections. The flux-recovery problem is slightly different in this

Table 2. Statistics on the recovered flux for each channel for the three approaches and for various flux bands (in mJy). The best results are displayed in bold.

Channel	Method	Flux < 1 Jy			1 Jy ≤ Flux < 5 Jy			Flux ≥ 5 Jy		
		Bias	RMSE	MAD	Bias	RMSE	MAD	Bias	RMSE	MAD
K	FIT-C	114.1	392.9	292.6	112.4	384.9	292.2	67.6	650.71	404.92
	FIT-L	114.8	380.9	286.0	96.0	471.3	298.7	90.0	518.5	317.9
	SPSR	41.5	334.8	260.4	47.1	362.6	272.6	-16.9	364.7	247.0
Ka	FIT-C	70.8	357.4	283.0	50.4	392.6	312.7	-26.9	548.1	356.3
	FIT-L	69.4	349.4	281.1	38.7	468.4	313.7	30.6	508.5	378.9
	SPSR	25.7	259.0	205.8	6.6	292.5	220.9	1.3	369.5	289.0
Q	FIT-C	36.8	319.9	259.0	29.5	356.0	284.8	-38.1	380.5	292.2
	FIT-L	38.5	332.2	258.9	10.8	400.8	290.0	13.0	310.10	249.2
	SPSR	11.0	217.3	171.4	10.2	254.6	184.7	21.9	242.8	192.8
V	FIT-C	-9.1	287.7	235.7	34.7	296.5	241.5	37.0	281.5	200.1
	FIT-L	-0.55	321.3	240.4	16.0	353.8	245.0	42.6	304.5	214.4
	SPSR	-1.41	263.9	208.9	13.8	265.2	208.2	1.8	231.3	167.5
W	FIT-C	-56.1	281.7	224.2	18.0	266.8	207.6	2.6	346.8	254.3
	FIT-L	-54.7	280.1	222.8	20.0	265.9	206.9	13.6	283.6	224.4
	SPSR	-62.0	296.5	229.6	30.4	273.8	208.6	11.7	257.6	202.4

case, since the catalogue is first subject to Eddington bias (excess detection in positive fluctuations of the background, which translates into higher biases in estimating the flux and impacts low-flux sources), and then the task is also to recover as accurately as possible the spectra of sources that can be lower than the detection level in some of the channels (e.g. in higher frequencies for radio sources).

3.3.1. Catalogue generation

As was previously done for the *WMAP* official point-source catalogue, we therefore constructed a catalogue by first applying a matched filter to each channel, taking into account second-order statistics of CMB and noise in spherical harmonic space (Bennett et al. 2012; Tegmark & de Oliveira-Costa 1998). Sources were considered detected if the value in the filtered map was higher than 5σ of the filtered background (computed locally). This resulted in 689 (442, 415, 248, and 172) sources for channel K (Ka, Q, V, and W). Compared with the results obtained in Table 1, a sub-population of the previous catalogue is essentially considered at this step: sources with sufficiently high local contrast to the diffuse background that they can be internally detected. After merging of the sources, the final simulation-derived multichannel catalogue contained 724 sources, and we processed these in the same way as described in the previous paragraph. However, after merging, the sources considered for a channel with this approach are no longer a sub-population of the previous catalogue: sources detected in other bands are considered as well.

3.3.2. Flux recovery and point-source subtraction

The same working conditions were set for SPSR as in the previous simulations. The statistics obtained as in Table 2 are presented in Table 3. The values reported for medium- or high-flux point sources in the first channels are similar if not identical, because this sub-population did not change from the previous catalogue. However, significant discrepancies from Table 2 can be noted. First, the biases have increased for all methods for low-flux sources, an indication for Eddington bias. Note, however, that the bias has increased less for SPSR than for FIT-C or FIT-L in absolute values, leading to even larger differences between the methods (SPSR leads to the lowest biases for all channels, with a bias now up to more than 100mJy lower for sources lower than 1Jy in channel K). This is also true for the RMSE of these low-flux sources, which essentially did not change for SPSR and was slightly degraded for FIT-C or FIT-L. Note that most of the low-flux sources in the catalogue are not detected in the W band, which might explain

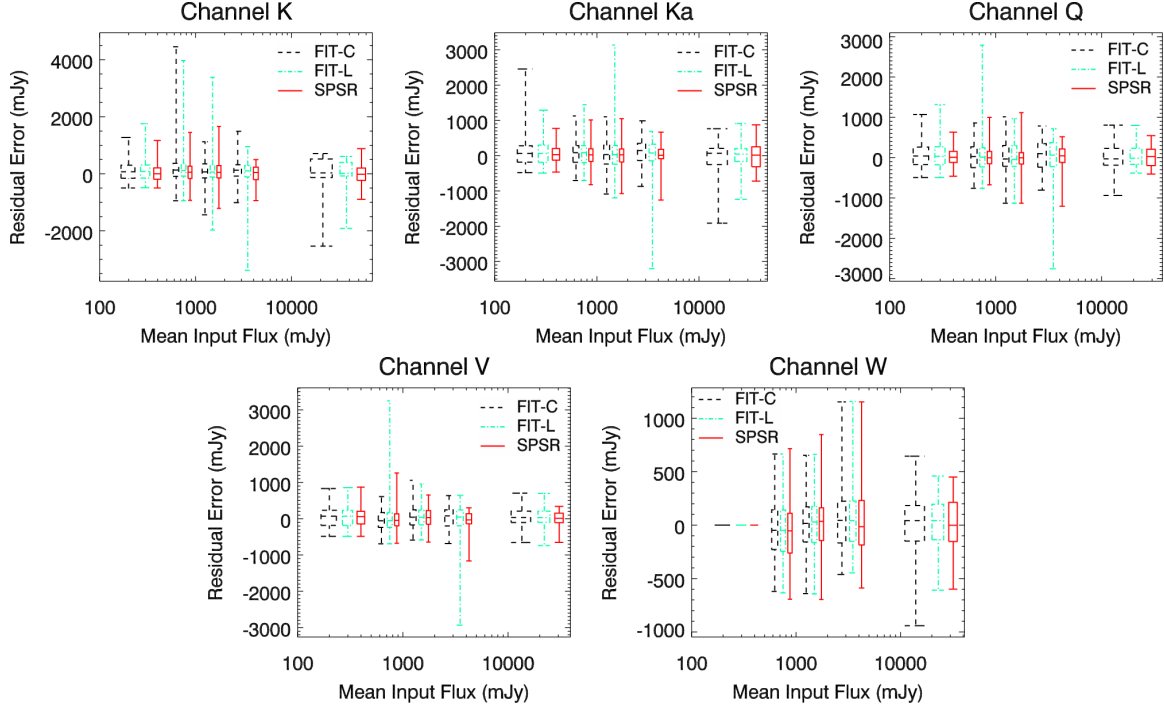


Figure 4. Statistics on error flux using the various approaches for the five channels for five different flux bands ($\text{flux} < 500 \text{ mJy}$, $500 \text{ mJy} \leq \text{flux} < 1 \text{ Jy}$, $1 \text{ mJy} \leq \text{flux} < 2 \text{ Jy}$, $2 \text{ Jy} \leq \text{flux} < 5 \text{ Jy}$, $\text{flux} \geq 5 \text{ Jy}$). Quartiles and extreme values are plotted.




Figure 5. Point-source residuals with noise for the various approaches in a projected region of channel Q, illustrating the lower bias obtained for this channel with the proposed approach. These maps were obtained by subtracting the diffuse component from the point-source-subtracted map. In the first two maps, three negative regions can be visually detected because of the overfitting of the point-source flux in the FIT-C and FIT-L approaches, which is different from the SPSR residual map. This does not display biases that are significantly higher than the noise level.

why SPSR gave better RMSE for medium-flux sources than FIT-C or FIT-L, while this was not the case in the noise-limited catalogue. The proposed approach in this multichannel scenario therefore leads to the lowest errors, even for flux estimated in channel W.

The distribution of errors is reported in Figure 7. SPSR seems as robust to large flux errors in the internal catalogue as in the noise-limited case with respect to FIT-C or FIT-L. Note the distribution of errors for channel W for low-flux sources (lower than 500mJy), which were indeed not detected in this




Figure 6. Projected regions with large discrepancy between the proposed approach and the local low-order polynomial background modelling. Images correspond to the simulated data, the diffuse component with noise, and the point-source-subtracted data with the three approaches. SPSR in this case allows a better subtraction of the point-source contribution because the local variations in the background are not captured by the other two approaches.

channel (compare with Figure 4). Large flux errors are observed for all methods, but are less pronounced for SPSR than for FIT-C or FIT-L.

As a summary for the internal catalogue, the most striking difference compared with the noise-limited catalogue occurs in low-flux sources that are now subject to Eddington bias, whereas for high-flux the results are essentially similar. Better relative performance in this situation in terms of bias and RMSE or MAD may be attributed to two different phenomena: the proposed approach would be less affected by Eddington bias and/or lead to better estimates for sources not detected in the band.

4. WMAP 9-year processing

The proposed approach was then applied to subtract point-source emission from *WMAP* 9-year data. The maps were processed as follows: the differential assemblies data deconvolved with the asymmetrical part of the beam were averaged to obtain 9-year frequency band data (Bennett et al. 2012). The beam considered for each channel was obtained as the mean over all axisymmetric beams provided for the differential assemblies at that frequency. A catalogue was then built by merging the two catalogues provided by the *WMAP* collaboration, resulting in 628 considered point sources.

The same processing (with the same working conditions) as for the simulations was then performed for both SPSR and the local low-order polynomial minimization. In the absence of ground truth, it is obviously difficult to quantitatively assess the relative performance of the algorithm as in the simulation. Only extreme cases that visually illustrate the performance of the approaches are therefore presented in Figs. 8 and 9.

Table 3. Statistics on the recovered flux for each channel using the multichannel internal catalogue for the three approaches and for various flux bands (in mJy). Again, the best results are displayed in bold.

Channel	Method	Flux < 1 Jy			1 Jy ≤ Flux < 5 Jy			Flux ≥ 5 Jy		
		Bias	RMSE	MAD	Bias	RMSE	MAD	Bias	RMSE	MAD
K	FIT-C	171.3	437.7	311.7	112.8	385.1	292.9	68.4	650.6	404.2
	FIT-L	168.5	422.6	305.3	97.5	472.6	300.1	90.0	518.5	317.5
	SPSR	70.3	337.6	257.7	46.6	364.3	274.1	0.4	380.1	275.0
Ka	FIT-C	115.5	381.0	296.0	50.4	392.6	312.7	-26.9	548.1	356.3
	FIT-L	115.5	371.7	296.7	38.7	468.4	313.7	30.6	508.5	378.9
	SPSR	49.3	268.2	210.7	6.6	293.9	221.9	11.9	376.0	289.2
Q	FIT-C	75.4	334.1	267.1	30.0	355.9	284.4	-26.2	373.7	280.3
	FIT-L	79.1	354.7	269.5	11.5	400.6	289.3	22.6	305.6	239.6
	SPSR	30.4	227.2	177.1	6.9	255.5	185.1	22.4	236.9	191.5
V	FIT-C	27.3	286.4	228.4	37.9	300.0	243.5	21.1	248.2	178.0
	FIT-L	32.8	314.8	230.8	16.3	353.7	245.3	53.5	312.8	224.0
	SPSR	17.6	259.6	202.8	7.1	274.0	211.8	-23.3	230.1	172.8
W	FIT-C	33.8	289.1	201.5	-11.0	357.2	230.0	4.4	346.1	254.3
	FIT-L	32.3	286.6	200.7	-10.7	353.8	227.2	13.7	285.6	225.4
	SPSR	30.2	277.7	210.3	10.3	302.8	220.3	-8.3	253.7	198.1

Figure 8 represents a patch centred on the source with the strongest difference between SPSR and FIT-L for channel V. Because the difference between the approaches is only a fraction of the background fluctuations at larger scales than the sources, visually no difference can be observed for that channel. However, Figure 9 illustrates an extreme case for channel K, suggesting some over-estimation of the flux in a complex background scenario with FIT-C and FIT-L, but not with SPSR, as reported in Figure 6 in the simulation.

5. Reproducible research

Product Name	Description
<i>WMAP</i> 9-year products:	
WMAP9_data_chanK_ptsub.fits	SPSR point-source-free <i>WMAP</i> 9-year K channel
WMAP9_data_chanKa_ptsub.fits	SPSR point-source-free <i>WMAP</i> 9-year Ka channel
WMAP9_data_chanQ_ptsub.fits	SPSR point-source-free <i>WMAP</i> 9-year Q channel
WMAP9_data_chanV_ptsub.fits	SPSR point-source-free <i>WMAP</i> 9-year V channel
WMAP9_data_chanW_ptsub.fits	SPSR point-source-free <i>WMAP</i> 9-year W channel
Software products (IDL):	
mrs_sparse_pointsource_removal.pro	removes point sources from a spherical map (requires HealPix and iSAP).
WMAP9_DATA_routines.pro	routines to obtain the <i>WMAP</i> 9-year maps and beams required by SPSR (requires HealPix and iSAP).
wmap9_remove_point_sources.pro	script that applies SPSR to all <i>WMAP</i> 9-year channels (requires HealPix and iSAP).

Table 4. List of products made available in this paper in the spirit of reproducible research, available at <http://www.cosmostat.org/Products>.

In the spirit of participating in reproducible research, we make all codes and resulting products that constitute the main results of this paper public. In Table 4 we list all products that are made freely available as a result of this paper, and which are available at <http://www.cosmostat.org/Products>.

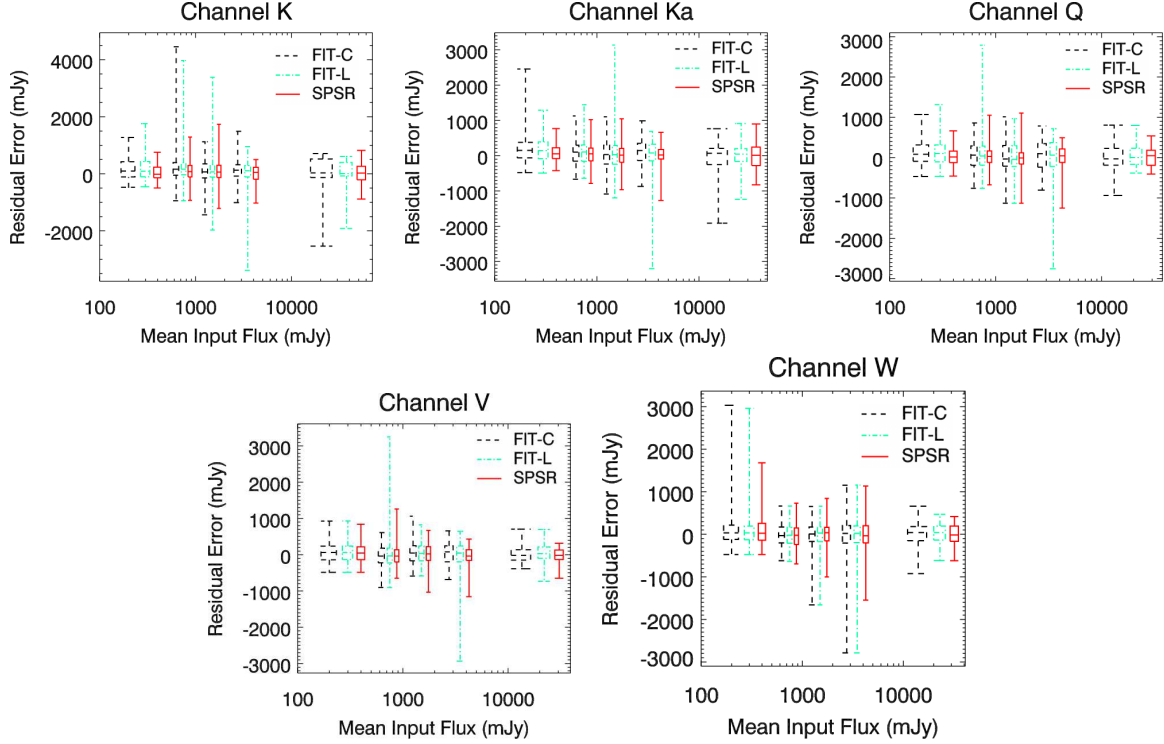


Figure 7. Statistics on error flux computed from internal catalogue and using the various approaches for the five channels for five different flux bands (flux < 500 mJy, $500 \text{ mJy} \leq \text{flux} < 1$ Jy, $1 \text{ mJy} \leq \text{flux} < 2$ Jy, $2 \text{ Jy} \leq \text{flux} < 5$ Jy, flux ≥ 5 Jy). Quartiles and extreme values are plotted.

For the five *WMAP* 9-year channels, SPSR was applied using the following command line in the open source package *iSAP* software¹:

```
> map_psfree =  
  mrs_sparse_pointsource_removal(Map, GalMask, BeamInfo, StdMap, Niter=Niter).
```

For Ka to W maps, Niter was fixed to 9750, and to 13350 for the K channel. Galmask is the *WMAP* point-source catalogue mask², StdMap is the map containing the standard deviation of the noise per pixel, and BeamInfo is a structure containing the *WMAP* 9-year beam at each point-source position. The code *WMAP9_DATA_routines.pro* can be used to obtain the required maps and beams and to call the *mrs_sparse_pointsource_removal* routine to derive the final products, as illustrated in the script *wmap9_remove_point_sources.pro*.

6. Conclusions

We proposed a new approach for detected point-source flux estimation and subtraction. Compared with the standard approach, which estimates point-source flux according to local low-order polynomial models of the background, the proposed technique is based on a global modelling of the background, which is assumed to be sparse in spherical harmonics to better capture its fluctuations. Bright point-source emissions decrease the sparsity of this background, which is the key phenomenon driving the separation process. An algorithm was adapted from recent convex optimization developments to solve the corresponding inverse problem.

We evaluated the proposed technique as well as techniques used in *WMAP* and *Planck* collaboration on realistic simulations of the *WMAP* microwave sky. In a noise-limited catalogue, except for channel W, where noise leads to faint differences between the estimates and a slightly poorer estimate for the proposed approach for low-flux sources (about 6%), our approach out-performs local polynomial fitting. In the internally derived catalogue, SPSR also consistently leads to

- the lowest biases in the first three channels (up to 100mJy lower bias for sources < 1 Jy in channel K)

¹ <http://jstarck.free.fr/isap.html>

² http://lambda.gsfc.nasa.gov/data/map/dr5/ancillary/masks/wmap_point_source_catalog_mask_r9_9yr_v5.fits

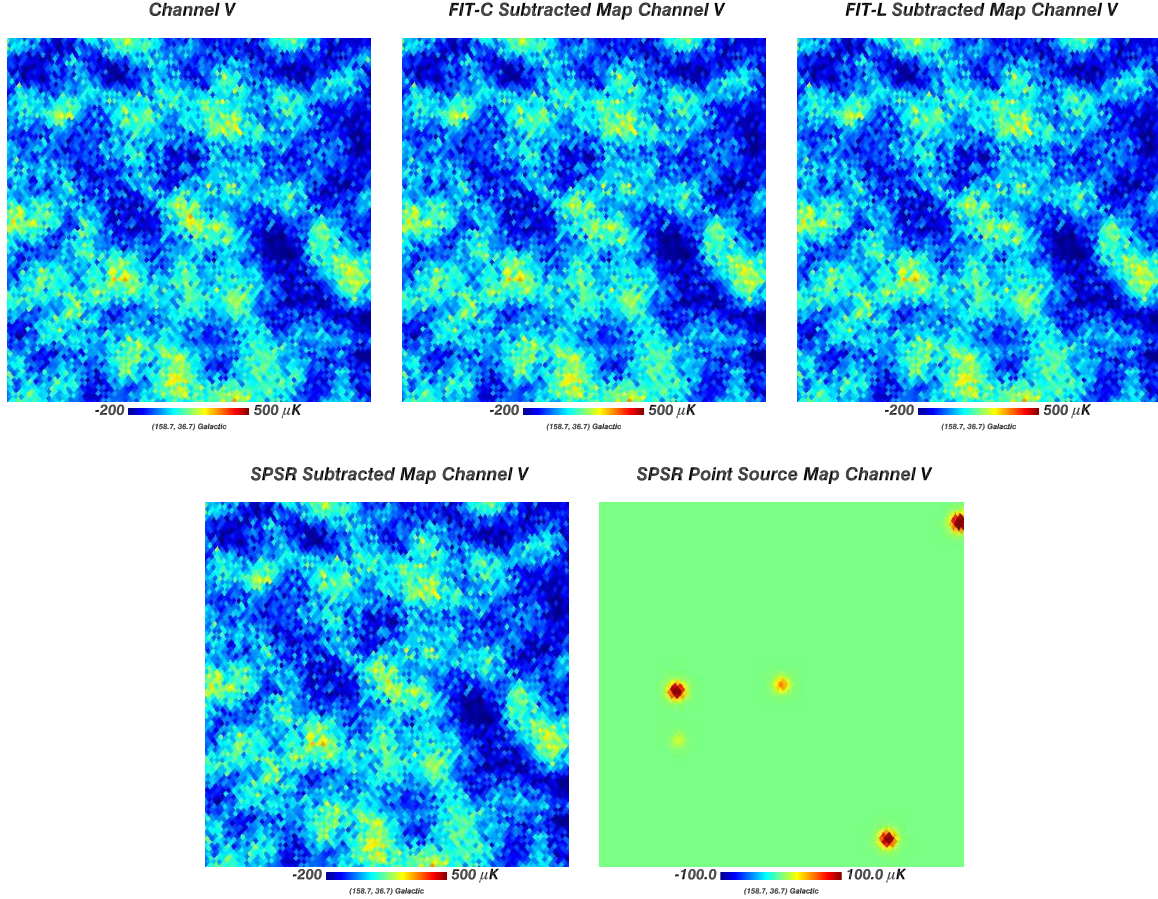


Figure 8. Projected region in channel V of the *WMAP* 9-year data showing the largest discrepancy between SPSR and FIT-C or FIT-L estimated fluxes for two fitted sources. Images correspond to the *WMAP* data, the point-source-subtracted data with FIT-C, FIT-L, and SPSR approaches, and point-source-estimated data using SPSR. The differences here between the approaches are a fraction of noise and the diffuse component levels.

- the lowest RMSE in all channels and all flux bands (with values decreased by at least 5%, 25%, 28%, 7%, and 3% and up to 35%, 28%, 32%, 9%, and 14% for channels K to W)
- more robust point-source subtraction in complex background, as illustrated in specific examples
- consequently better multichannel estimates of the point-source fluxes, which is useful to estimate more accurate spectra.

This technique was finally applied to the *WMAP* 9-year data deconvolved with the beam asymmetries and the resulting point-source-subtracted maps are available on-line. We focused in this paper on point sources. Even if the proposed model for extended compact sources is rather simple in SPSR, it is sufficient to derive a background accurate enough to improve on point-source flux recovery compared with FIT-C or FIT-L. However, better modelling of the extended compact sources is still required if they need be subtracted from the data. We left this task for future work.

Acknowledgment

This work was funded by the European Research Council grant SparseAstro (ERC-228261), and the Swiss National Science Foundation (SNSF). We used the Healpix³ software (Górski et al. 2005), the iSAP¹ software, and *WMAP* data⁴.

³ <http://healpix.sourceforge.net>

⁴ <http://map.gsfc.nasa.gov>

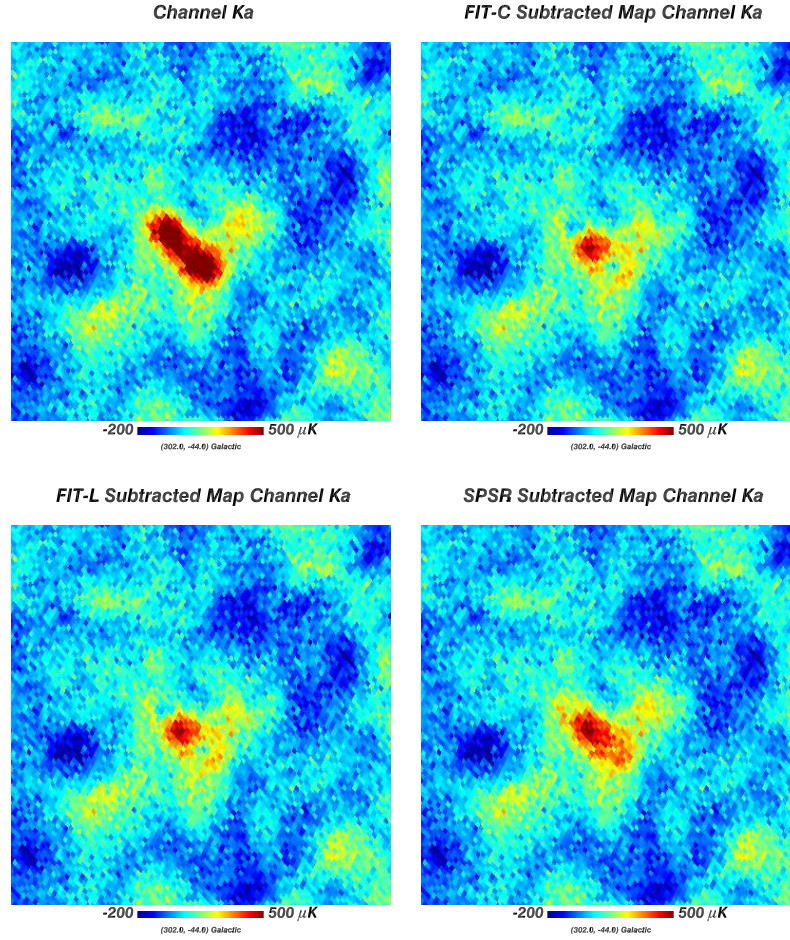


Figure 9. Projected region in *WMAP* 9-year data showing the large discrepancy in between SPSR and FIT-C or FIT-L estimated fluxes for two fitted sources. Images correspond to the *WMAP* data, the point-source-subtracted data with FIT-C, FIT-L, and SPSR approaches. The positive ring around two the fitted regions as well as a negative peak inside these regions seem to indicate that FIT-C or FIT-L are here over-fitting. The same phenomenon is absent for SPSR.

References

- Abrial, P., Moudden, Y., Starck, J., et al. 2007, *Journal of Fourier Analysis and Applications*, 13, 729
- Abrial, P., Moudden, Y., Starck, J., et al. 2008, *Statistical Methodology*, 5, 289
- Argüeso, F., Salerno, E., Herranz, D., et al. 2011, *Monthly Notices of the Royal Astronomical Society*, 414, 410
- Barreiro, R. B., Sanz, J. L., Herranz, D., & Martínez-González, E. 2003, *Monthly Notices of the Royal Astronomical Society*, 342, 119
- Becker, S. R., Candès, E. J., & Grant, M. C. 2011, *Mathematical Programming Computation*, 3, 165
- Bennett, C. et al. 2012, *ArXiv preprint*, 1212.5225
- Bertin, E. & Arnouts, S. 1996, *Astronomy and Astrophysics, Supplement Series*, 117, 393
- Briceño Arias, L. & Combettes, P. 2011, *SIAM Journal on Optimization*, 21, 1230
- Bucher, M. & Louis, T. 2012, *Monthly Notices of the Royal Astronomical Society*, 424, 1694
- Carvalho, P., Rocha, G., & Hobson, M. P. 2009, *Monthly Notices of the Royal Astronomical Society*, 393, 681
- Cayón, L., Sanz, J. L., Barreiro, R. B., et al. 2000, *Monthly Notices of the Royal Astronomical Society*, 315, 757
- Chambolle, A. & Pock, T. 2011, *Journal of Mathematical Imaging and Vision*, 40, 120
- Chen, S. & Donoho, D. 1998, in *IEEE International Conference on Acoustics, Speech and Signal Processing*, Vol. 3, 1865–1868
- Chen, X. & Wright, E. L. 2009, *The Astrophysical Journal*, 694, 222
- Combettes, P. L. & Pesquet, J.-C. 2012, *Set-Valued and Variational Analysis*, 20, 307
- Delabrouille, J., Betoule, M., Melin, J.-B., et al. 2013, *Astronomy and Astrophysics*, 553, A96
- Donoho, D. & Elad, M. 2003, *Proceedings of the National Academy of Sciences*, 100, 2197
- Donoho, D. & Huo, X. 2001, *IEEE Transactions on Information Theory*, 47, 2845–2862
- Dupé, F.-X., Rassat, A., Starck, J.-L., & Fadili, M. J. 2011, *Astronomy and Astrophysics, Supplement Series*, 534, A51
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *Astrophysical Journal*, 622, 759

- Guglielmetti, F., Fischer, R., & Dose, V. 2009, *Monthly Notices of the Royal Astronomical Society*, 396, 165
- Herranz, D. & Sanz, J. L. 2008, *IEEE Journal of Selected Topics in Signal Processing*, 2, 727
- Hobson, M. P. & McLachlan, C. 2003, *Monthly Notices of the Royal Astronomical Society*, 338, 765
- Kim, J., Naselsky, P., & Mandolesi, N. 2012, *The Astrophysical Journal Letters*, 750, L9
- Lanz, L., Herranz, D., Lopez-Caniego, M., et al. 2011, ArXiv preprint, 1110.6877
- Larson, D., Dunkley, J., Hinshaw, G., et al. 2011, *The Astrophysical Journal Supplement Series*, 192, 16
- López-Caniego, M., Herranz, D., González-Nuevo, J., et al. 2006, *Monthly Notices of the Royal Astronomical Society*, 370, 2047
- Martínez-González, E., Diego, J. M., Vielva, P., & Silk, J. 2003, *Monthly Notices of the Royal Astronomical Society*, 345, 1101
- Nolta, M. R., Dunkley, J., Hill, R. S., et al. 2009, *The Astrophysical Journal Supplement Series*, 180, 296
- Paykari, P. & Starck, J.-L. S. 2012, *Cosmic Microwave Background Data Analysis*, ed. M. J. Way, J. D. Scargle, K. M. Ali, & A. N. Srivastava, 55–87
- Perotto, L., Bobin, J., Plaszczyński, S., Starck, J., & Lavabre, A. 2010, *Astronomy and Astrophysics*, 519, A4+
- Planck Collaboration. 2011a, *Astronomy and Astrophysics*, 536, A7
- Planck Collaboration. 2011b, *Astronomy and Astrophysics*, 536, A14
- Planck Collaboration. 2013a, ArXiv e-prints, 1303.5072
- Planck Collaboration. 2013b, ArXiv e-prints, 1303.5088
- Ramos, E. P. R. G., Vio, R., & Andreani, P. 2011, *Astronomy and Astrophysics*, 528, A75
- Savage, R. S. & Oliver, S. 2007, *Astrophysical Journal*, 661, 1339
- Scodeller, S. & Hansen, F. K. 2012, *Astrophysical Journal*, 761, 119
- Scodeller, S., Hansen, F. K., & Marinucci, D. 2012, *Astrophysical Journal*, 753, 27
- Starck, J.-L., Donoho, D. L., Fadili, M. J., & Rassat, A. 2013a, *Astronomy and Astrophysics*, 552, A133
- Starck, J.-L., Elad, M., & Donoho, D. 2004, *Advances in Imaging and Electron Physics*, 132, 287
- Starck, J.-L., Elad, M., & Donoho, D. 2005, *IEEE Transactions on Image Processing*, 14, 1570
- Starck, J.-L., Fadili, M. J., & Rassat, A. 2013b, *Astronomy and Astrophysics*, 550, A15
- Starck, J.-L., Moudden, Y., Abrial, P., & Nguyen, M. 2006, *Astronomy and Astrophysics*, 446, 1191–1204
- Tegmark, M. & de Oliveira-Costa, A. 1998, *Astrophysical Journal Letters*, 500, L83
- Vielva, P., Martínez-González, E., Cayón, L., et al. 2001, *Monthly Notices of the Royal Astronomical Society*, 326, 181
- Vielva, P., Martínez-González, E., Gallegos, J. E., Toffolatti, L., & Sanz, J. L. 2003, *Monthly Notices of the Royal Astronomical Society*, 344, 89
- Vio, R., Andreani, P., Ramos, E. P. R. G., & da Silva, A. 2012, ArXiv e-prints, 1206.4536
- Wright, E. L., Chen, X., Odegard, N., et al. 2009, *Astrophysical Journal, Supplement Series*, 180, 283
- Zacchei, A., Maino, D., Baccigalupi, C., et al. 2011, *Astronomy and Astrophysics*, 536, A5